Acta Astronautica 103 (2014) 322–332

Contents lists available at ScienceDirect

Acta Astronautica

journal homepage: www.elsevier.com/locate/actaastro

Closed loop terminal guidance navigation for a kinetic impactor spacecraft

Shyam Bhaskaran*, Brian Kennedy

Outer Planet Navigation Group, Mission Design and Navigation Section, Jet Propulsion Laboratory, California Institute of Technology, MS 264-820, 4800 Oak Grove Dr., Pasadena, CA 91109, United States

ARTICLE INFO

Article history:
Received 30 October 2013
Received in revised form
11 February 2014
Accepted 25 February 2014
Available online 5 March 2014

Keywords:
Kinetic impact
Asteroid deflection
Autonomous navigation
Closed loop guidance

ABSTRACT

A kinetic impactor spacecraft is a viable method to deflect an asteroid which poses a threat to the Earth. The technology to perform such a deflection has been demonstrated by the Deep Impact (DI) mission, which successfully collided with comet Tempel 1 in July 2005 using an onboard autonomous navigation system, called AutoNav, for the terminal phase of the mission. In this paper, we evaluate the ability of AutoNav to impact a wider range of scenarios that a deflection mission could encounter, varying parameters such as the approach velocity, phase angle, size of the asteroid, and the attitude determination accuracy. In particular, we evaluated the capability of AutoNav to impact 100–300 m size asteroids at speeds between 7.5 and 20 km/s at various phase angles. Using realistic Monte Carlo simulations, we tabulated the probability of success of the deflection as a function of these parameters and find the highest sensitivity to be due to the spacecraft attitude determination error. In addition, we also specifically analyzed the impact probability for a proposed mission (called ISIS) which would send an impactor to the asteroid 1999RQ36. We conclude with some recommendations for future work.

1. Introduction

On July 4, 2005, the Deep Impact (DI) spacecraft successfully impacted the comet Tempel 1 using a 400 kg impactor while the mother ship flew by the comet at a distance of 500 km and captured images of the impact. This event marked the first hypervelocity impact of a small solar system body, and, although not its primary goal, the mission demonstrated that such an impact could be accomplished with present day technologies, and on a relatively modest budget. Beyond its science benefits, the mission is notable because the same technology could be used to someday save the Earth from a devastating impact from a Near Earth Asteroid (NEA).

The potential threat from NEAs has been documented extensively [1] and cannot be understated. There are several options for mitigating this threat, including destroying the asteroid using a nuclear device, or altering its trajectory via gravity tractoring, impacting with another spacecraft (kinetic energy deflection), or a combination of any of these methods. The kinetic energy deflection technique is the most straightforward and easiest to implement using currently available technologies. The technique relies on both the direct momentum transfer from the impact, plus whatever is added by the momentum of the resulting ejecta. The momentum enhancement factor, however, is unknown, and it would be advantageous to narrow down the ranges via a deflection experiment. Ideally, this experiment would have two components: an impacting spacecraft and a second spacecraft in the proximity of the asteroid which measures the deflection (see Ref. [2] for a discussion on one proposed experiment). The main difficulty is the terminal guidance navigation; precisely hitting the target at high velocities where there is

* Corresponding author. Tel.: +1 818 354 3152; fax: +1 818 393 7413.
E-mail address: Shyam.Bhaskaran@jpl.nasa.gov (S. Bhaskaran).

http://dx.doi.org/10.1016/j.actaastro.2014.02.024
0094-5765

little time to react and must be done largely autonomously onboard the spacecraft is a challenge. However, DI has demonstrated that this is feasible, and the next step is to show how the point solution used to successfully impact Tempel 1 can be generalized to cover the range of scenarios for an asteroid deflection.

DIs impact was made possible by the onboard closed loop autonomous navigation system (called AutoNav). The filter settings and sequence of events performed by AutoNav to achieve the impact were determined through simulations to maximize the probability of impact for this particular spacecraft and scenario. Some of the critical parameters were the size of the comet (roughly 6 km in diameter), approach velocity (10.5 km/s), and the approach phase angle (62 deg). The parameters for a NEA deflection could be considerably different. The range of possible sizes could be as low as 100 m in diameter or less, the approach velocity could range from below 10 km/s to as high as 20 km/s, and the approach phase could range from near 0 deg to near 180 deg (i.e., fully lit or no solar illumination). Thus, before a deflection mission is undertaken, it is important to understand how the terminal guidance will perform under a range of conditions.

In this paper, we expand the experience base of using AutoNav for terminal guidance in an asteroid deflection mission and parameterize the probability of achieving a successful impact on a sample set of deflection missions. We first determined the ranges of parameters for deflection missions through example scenarios from the literature in order to define the scope of conditions that AutoNav needs to handle. This is combined with information on several important spacecraft hardware considerations that affect AutoNav performance and their interaction with the mission parameters. Finally, Monte Carlo simulations of AutoNav terminal guidance were performed on selected scenarios, varying all relevant parameters, to obtain statistics on probabilities of impact. In addition, a specific case, the ISIS mission, was analyzed. This mission is a proposal to impact the asteroid 1999RQ36 in early 2021 while the OSIRIS-REx mission is already there, such that the latter can observe the impact and measure the subsequent deflection. The simulations are run using high fidelity models which describe the spacecraft trajectory, spacecraft attitude errors, and ephemeris errors of the target body. The simulations include: (1) generation of realistic images using triaxial ellipsoid shape model and the parameters of the camera, (2) determination of the orbit using a batch least-squares filter, and (3) maneuver targeting to achieve impact conditions.

2. Deep space navigation

Before we describe the AutoNav system used for terminal guidance, we provide here a brief overview of deep space navigation in general. The first step in any deep space mission is to design the reference trajectory to achieve the target conditions, which, in this case, is to impact a candidate set of asteroids or a particular one. The techniques to design these trajectories, optimizing parameters of interest such as fuel expenditure or time-of-flight, are covered elsewhere (see Ref. [3] for an example) and not in the scope of this paper. However, the ensemble set of trajectories found in Ref. [3] was used to define the scope of the problem and will be referred to later.

Once a particular set of trajectories is found, the next step is to analyze them for its flyability from a navigation perspective. This includes performing linear covariance analyses to determine the navigation delivery performance and statistics on maneuver ΔV requirements for the mission. The covariance analysis involves using a realistic tracking schedule to simulate tracking data and fitting the data in a least-squares process to determine the orbit. The tracking data includes two-way Doppler and range, plus optical data (images of the target object taken with an onboard camera). The least-squares fitting provides error covariances which statistically describe how accurate the spacecraft can be navigated. This also provides inputs to simulate the maneuvers required to achieve the target; this is done via Monte Carlo simulations which sample the orbit determination errors to design the maneuvers. The end result describes the amount of fuel needed to deliver the vehicle to the target, or in our case, to the start of the autonomous phase. When the mission is actually flown, real tracking data replaces the simulated ones, and the best fit orbit is used to design the maneuvers. If all has gone according to plan, the delivery errors to the target and the fuel required will be within the statistics of the pre-flight analysis. We will not cover the details of navigation for the launch, cruise, and target approach mission phases; for a general overview deep space navigation techniques, see Refs. [4,5] show an example of its application for a comet flyby mission which closely resembles that of an asteroid impactor.

For this paper, we are concerned with the impactor's terminal guidance navigation, which we define as the phase beginning roughly 2–3 h prior to impact. At this stage, ground navigation techniques are impractical due to the round-trip light time and the onboard navigation system, AutoNav, must be used. Details of the AutoNav system and how it was used on various missions can be found in Refs. [6–8]; here we will briefly describe the orbit determination filter as used by AutoNav in a subsequent section. First we describe some other considerations necessary to completely describe the problem set up.

3. Asteroid ephemeris

In order to successfully hit an asteroid with a spacecraft, two critical pieces of information are needed. The first is knowledge of the spacecraft's trajectory as described above. The second is knowledge of the target asteroid's orbit. Preliminary estimates of the latter are provided through ground-based observations of the asteroid, primarily from optical telescopes, but also in a limited number of cases from radar bounces off the asteroid [2]. The accuracy of this method is dependent on many factors including the density, quality, resolution, and geometry of the observations, the orbital characteristics of the asteroid, and the length of time from the last observation to the time of spacecraft approach and impact. In general, however, we can say that the accuracy of the orbit from ground-based observations as the spacecraft approaches will be in the tens of km range.

This is obviously not good enough to hit an object which may only be a 100 m across, and thus eventually the ephemeris knowledge must be reduced to better than this level. The key point, however, is that it is not the asteroid's (or the spacecraft's for that matter) heliocentric orbit that must be known accurately but the relative orbit of the spacecraft to the asteroid. This information is obtained by the onboard camera, which will now be described.

4. Camera

Optical navigation, or OpNav, is the science of using an onboard camera as a navigation device [9]. OpNav involves taking images of solar system bodies against a star background; precisely locating the center of the body provides an angular measurement than can be used as data for an orbit determination filter. Unlike Doppler and ranging which provides an Earth-relative measurement, OpNav data type is the only one that provides navigation information relative to a target body. Because it is an angular measure, the data becomes stronger as the spacecraft gets closer to the target.

In recent times, the camera used for OpNav is usually a combination of camera lenses focusing light onto a Charge-Coupled-Device (CCD) array. The light energy from a source in the form of photons is measured in each pixel of the array and output, through an analog to digital converter, as a "data number (DN)". The signal-to-noise (S/N) ratio of the light source is a function of various properties of the camera optics and electronics; in order for a signal to be usable, the ratio must be greater than around 3. The resolution and field-of-view (FOV) of the camera are determined by a combination of the camera lens focal length, and the size and number of the pixel array on the CCD. One simple measure of the resolution is the IFOV, the angular resolution of a single pixel, which can be computed as

$$IFOV = 1/FK, \tag{1}$$

where F is the focal length in mm, and K is the reciprocal of the pixel dimensions, expressed as pixels/mm. For example, the medium resolution camera on DI had a focal length of 2101 mm and a value of K of 47.6 pixels/mm, giving an effective IFOV of 10 μrad. This parameter is particularly important for the impactor application because the IFOV determines the accuracy of an OpNav measurement as well as when an asteroid becomes "resolved", that is, when its angular extent in the camera FOV is larger than a pixel. Even when the asteroid is unresolved, however, its light is spread over several pixels due to diffraction, and to a lesser extent, slight defocusing of the lens. This spread is defined by the camera's pointspread function (PSF), which can be modeled as a Gaussian function, and typical values of the PSF disperse the light from the unresolved source such that the central peak contains 30–40% of the total light from the source. Note that stars in the FOV are also unresolved and look very similar to unresolved objects.

Detection of the target asteroid depends on many factors including the sensitivity and noise of the camera, the size, shape, orientation, approach phase angle, exposure duration, and the stability of the spacecraft over the exposure duration. Because the asteroids in consideration for an impactor are small (generally less than 300 m in diameter), the detection will usually be fairly late on approach. The exact time will depend on the specific scenario, but overall. It is expected that the S/N will not be sufficient for detection until a day or two before impact at the earliest. If detection does occur over a day before impact, then the initial target relative orbit determination can be done on the ground, and a maneuver uplinked to remove the tens of km of asteroid ephemeris error described earlier. If detection is within a day or less, then the turnaround time for ground in the loop processing becomes problematic and the entire targeting process must be done using AutoNav. We define the terminal guidance phase as the part where the onboard process takes over.

Once detection occurs, the center of the object in the camera FOV must be precisely determined, a process called centerfinding. For stars and unresolved asteroids, several methods are available which fit a Gaussian or other function to the diffused signal [9]. Because of the PSF, this methodology can locate the center of a source to much better than a pixel. For stars whose cataloged positions are well known and have good S/N, the accuracy of the centerfinding can typically be as low as 0.02 pixels. Unresolved targets are typically determined to the 0.1–0.2 pixel level. As the asteroid becomes resolved, the Gaussian model no longer resembles the image as the shape of the asteroid becomes observable. In this case, the center can be found using a simple brightness moment algorithm. The accuracy of this method to locate the center of figure of the asteroid varies depending on its shape and the approach phase angle (the asteroid-centered angle between the incoming asymptote of the spacecraft and the sun), and will be proportional to the size of the asteroid in the camera FOV.

The primary purpose of locating stars in the image is to provide the inertial frame in which to measure the angle of the target asteroid relative to the spacecraft. Any error in this angle will be interpreted by the orbit determination filter as a translational motion of the spacecraft relative to the target. Normally, the inertial camera frame is given by the spacecraft's attitude control system (ACS), but as will be described in the next section, the accuracy of this information may not be sufficient for precise targeting. Thus, if available, visible stars in the image can be used to determine the camera attitude to well below the noise level of the angular measure of the target object. In principle, two stars are sufficient, but ideally three or more are needed to reduce catalog and centroiding errors. The question then shifts to the likelihood that at least two stars will be visible in an image taken of the target body. This is not easily answerable on a general level; it depends on several factors, including the brightnesses of the background stars in the direction of the approach asymptote, the sensitivity and dynamic range of the camera, and duration of the exposure. As will be shown in the simulations, the attitude knowledge is the single biggest factor in determining the likelihood of a successful impact. Thus, for a given target and approach trajectory, an important trade

is to determine whether the camera has the sufficient sensitivity and dynamic range to capture both the background stars and asteroid in the same image with sufficient S/N and below the CCD's saturation level, noting that the asteroid brightness will vary considerably as the asteroid gets closer. If the camera limitations prevent it, then alternate schemes can be worked where alternate frames are taken, bracketing a frame optimized to capture the asteroid by frames capturing the stars. In this manner, the attitude can be interpolated between the two star frames to get reasonable accuracy.

5. Spacecraft attitude control system

Knowledge and control of the spacecraft's orientation is done by the onboard ACS. To get accurate inertial attitude, star trackers are used; these relatively wide FOV cameras (typically 20–40 deg) match the pattern of stars seen in the image to cataloged star locations and then solve for the boresight pointing. The mounting of the star cameras is known, so the boresight pointing can be related to the defined Cartesian axes of the spacecraft. Generally, at least two star trackers, pointing in different directions, are used. This provides redundancy as well as better geometric information to improve attitude knowledge. Gyroscopic Inertial Measurement Units (IMUs) are also used to precisely propagate the attitude knowledge in the absence of star tracker information or the presence of tracker noise. The gyros provide attitude rate information, so they must be initialized with absolute information from star trackers, after which the rates are propagated. A Kalman filter is used to combine star tracker and IMU data to obtain the current attitude.

If stars are not available in the navigation camera to get a high accuracy pointing solution, then the orbit determination filter must rely on the attitude knowledge available from ACS. Because the star trackers have much wider FOVs than the navigation camera, its resolution, and hence accuracy, is coarser. Furthermore, the propagation of the attitude between star tracker measurements with the IMUs also suffer from gyro error sources. The interaction between the ACS attitude knowledge and the AutoNav orbit solution is a complex subject, and much of the analysis of AutoNav performance is dedicated to understanding this interaction (see Refs. [8,10] for examples). To summarize, we have found that the optimal scenario is to turn off star tracker updates prior to the initiation of AutoNav. This avoids the “star tracker spatial error” problem – a phenomenon where the attitude knowledge makes sudden and unexpected shifts as stars near the edge of the tracker frame go in and out of the frame. When AutoNav is started, there is an initial knowledge error of the attitude, then the remaining error is strictly due to gyro errors as the attitude is propagated. These gyro errors can be quantified statistically based on known performance specifications, primarily the gyro rate error (given in deg/h), and the angle random walk error (given in deg/$\sqrt{\text{h}}$). In the simulations described later, we use these specifications to see their effect on the terminal navigation. We used the terms “stellar mode” and “IMU mode” to refer to attitude knowledge obtained purely by star trackers and purely by IMUs, respectively.

6. Orbit determination

Orbit determination (OD) is the process of combining the OpNav measurements of the target location in a filter to estimate the full cartesian state of the spacecraft. We use a batch filter such that all the measurements are processed to get an estimate of the state at an epoch time. This state can then be propagated forward to get the state at other times. The measurements are the coordinates of the target brightness centroid in the camera frame; the horizontal coordinate is referred to as pixel, and the vertical is line. The relationship between the pixel/line coordinate and the line-of-sight direction to the asteroid is described in detail in Ref. [9] and is a function of the spacecraft's orientation and target relative state, and the camera parameters.

In the flight AutoNav code, the dynamic model uses full numerical integration of the equations of motion, with gravitational acceleration from the Sun and planets, plus solar radiation pressure acceleration, included in the force model. For all practical purposes, however, the dynamics of a spacecraft flying by a small asteroid can be modeled as straight line motion relative to the asteroid; solar radiation force has little effect in the 2–3 h on approach. Thus, for simplicity and computational speed, we model the spacecraft trajectory as a straight line with the only perturbing force coming from maneuvers executed by the spacecraft.

The only other effect on the measurement is the ACS errors in the case where stars are not visible in the navigation camera. Once again, for simplicity and without loss of generality, we can assume that the camera boresight is located along one of the spacecraft axes. Then, attitude errors in the axes perpendicular to the boresight are mapped directly into pixel and line camera coordinates. Attitude errors around the boresight will have little effect on the observed location of the target given the magnitudes of the errors we are considering. We then model the ACS error as a bias and rate in both pixel and line directions.

To complete the batch filter formulation [11], we describe the state vector to be estimated as

$$X = [x\ y\ z\ \dot{x}\ \dot{y}\ \dot{z}\ P\dot{P}\ L\ \dot{L}]. \tag{2}$$

where

$x, y, z =$ the asteroid − centered Cartesian position
$\times$vector of the spacecraft at the epoch time,
$\dot{x}, \dot{y}, \dot{z} =$ the asteroid − centered Cartesian velocity
$\times$vector of the spacecraft at the epoch time,
$P, \dot{P} =$ the attitude bias and rate in the
$\times$camera pixel coordinate,
$L, \dot{L} =$ the attitude bias and rate in the
$\times$camera line coordinate.

Define the dynamic parameters, $S = [x\ y\ z\ \dot{x}\ \dot{y}\ \dot{z}]$, and the observational bias parameters, $Q = [P\ \dot{P}\ L\ \dot{L}]$. The sensitivity matrix, H, is the partial derivative of the pixel/line centroid coordinates (p and l), with respect to the state

vector:

$$\mathbf{H} = \begin{bmatrix} \frac{\partial p}{\partial S} & \frac{\partial p}{\partial Q} \\ \frac{\partial l}{\partial S} & \frac{\partial l}{\partial Q} \end{bmatrix}. \tag{3}$$

The partials with respect to the dynamic parameters, S, are given in Ref. [9]; the partials with respect to the bias and rate parameters are simply:

$$\begin{bmatrix} \frac{\partial p}{\partial Q} \\ \frac{\partial l}{\partial Q} \end{bmatrix} = \begin{bmatrix} 1 & t & 0 & 0 \\ 0 & 0 & 1 & t \end{bmatrix}. \tag{4}$$

Given an initial guess at the initial state, an $nx2$ matrix of residuals, Y, can be computed by differencing the actual observed predicted pixel/line location of the asteroid with predicted values based on the initial guess. The least squares fit to get the best estimate of the state at the epoch, $\hat{x}$, is then

$$\hat{x} = (\mathbf{H}^T\mathbf{W}\mathbf{H})^{-1}\mathbf{H}^T\mathbf{W}\mathbf{Y}, \tag{5}$$

W is a 2×2 matrix of data weights, defined as

$$\mathbf{W} = \begin{bmatrix} 1/\sigma_p^2 & 0 \\ 0 & 1/\sigma_l^2 \end{bmatrix}, \tag{6}$$

where σ_p and σ_l are uncertainties on each pixel and line observation, respectively. The fitting process usually takes several iterations to converge on a solution, and given the epoch state solution, the future state at any time can be obtained by forward propagation.

Of particular interest is the miss distance at encounter, which ideally is 0 or some small value which is known to be on the surface of the target. A convenient way to represent this miss is the B-plane coordinate system. This frame is centered on the asteroid and perpendicular to the incoming asymptote (Fig. 1). The B-vector describes where the asymptote pierces the plane, the horizontal and vertical axes are denoted as T and R, and the coordinates of the piercing point are $B \bullet T$ and $B \bullet R$ (the projection of the B-vector on the T and R axes, respectively. The third, out-of-plane component is the Linearized Time-of-Flight, or LTOF. For a linear flyby, the B-plane can be computed from the cartesian via the following steps. Let V_∞ be the asteroid-relative incoming velocity vector of the spacecraft, and X is its position anywhere along

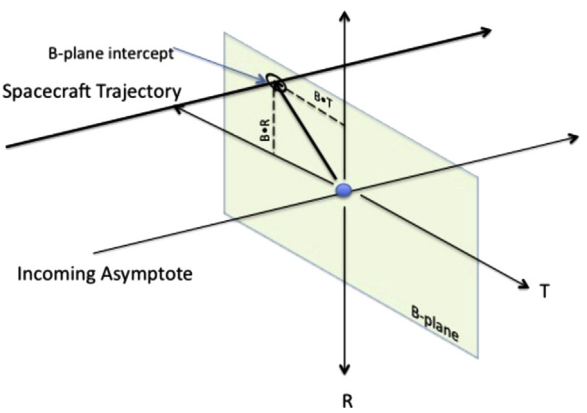

Fig. 1. B-plane coordinate system.

the trajectory. Then,

$$S = V_\infty / \| V_\infty \|, \tag{7}$$

$$C = \sqrt{1 - S_3^2}. \tag{8}$$

Then, the elements of T are computed as

$$T_1 = S_2 / C, \tag{9}$$

$$T_2 = -S_1 / C, \tag{10}$$

$$T_3 = 0 \tag{11}$$

The third axes, $R = S \times T$. The elements of the B-plane are then:

$$B \bullet R = X \bullet R^T, \tag{12}$$

$$B \bullet T = X \bullet T^T, \tag{13}$$

$$LTOF = X \bullet S^T / \| V_\infty \|. \tag{14}$$

7. Maneuver computation

Once the orbit is determined and the miss distance at the encounter determined, maneuvers must be computed to re-target the trajectory back to the intended aimpoint. The targeting algorithm is as follows. If we define Z to be a vector of targeted parameters in position, where

$$Z = [B \bullet R_t \;\; B \bullet T_t \;\; LTOF_t]^T, \tag{15}$$

denotes the desired B-plane state at impact (nominally all 0 if we wish to target the exact center of the body), and ΔV are the three components of the velocity adjustment used to control the trajectory, where

$$\Delta V = [\Delta V_x \;\; \Delta V_y \;\; \Delta V_z]^T, \tag{16}$$

then the sensitivity matrix $\mathbf{K}$ of the control parameters to the target is

$$\mathbf{K} = \begin{bmatrix} \frac{\partial B \bullet R_t}{\partial \Delta V_x} & \frac{\partial B \bullet R_t}{\partial \Delta V_y} & \frac{\partial B \bullet R_t}{\partial \Delta V_z} \\ \frac{\partial B \bullet T_t}{\partial \Delta V_x} & \frac{\partial B \bullet T_t}{\partial \Delta V_y} & \frac{\partial B \bullet T_t}{\partial \Delta V_z} \\ \frac{\partial LTOF_t}{\partial \Delta V_x} & \frac{\partial LTOF_t}{\partial \Delta V_y} & \frac{\partial LTOF_t}{\partial \Delta V_z} \end{bmatrix}. \tag{17}$$

Thus

$$K \Delta V = Z \tag{18}$$

and solving for the maneuver,

$$\Delta V = K^{-1} Z. \tag{19}$$

Starting with an initial reference trajectory and guess at the ΔV, the sensitivity matrix in Eq. (17) can be evaluated numerically by finite differencing methods. The initial guess is usually 0 which works sufficiently well in the dynamic regime we are working in. Several iterations are necessary to converge, i.e., when the miss distance (distance between the desired and achieved target) is less than some small tolerance. Normally, a single maneuver along the approach will not be enough as OD and maneuver execution errors will dominate initially; several maneuvers are used to fine tune the targeting. The exact number of maneuvers and their placement involve

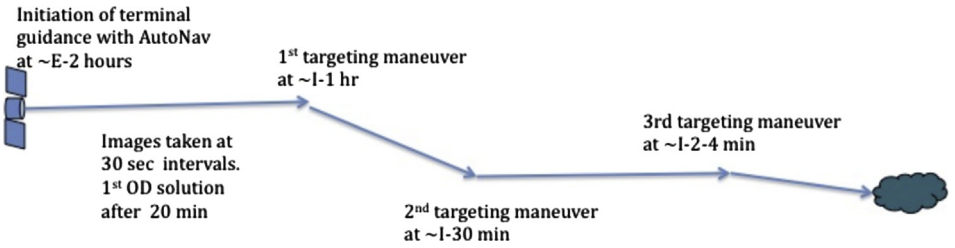

Fig. 2. Terminal guidance scenario.

a trade-off between the required accuracies, robustness to missed maneuvers, the amount of ΔV, and other factors and is not easily quantifiable. Thus, the end result is a solution based on criteria not necessarily related to the physics of the problem, and so for simplicity, we followed the DI approach as described in the next section. Also following the DI nomenclature, we refer to each maneuver as Impactor Targeting Maneuvers, or ITMs.

8. Terminal guidance scenario

The approach we take for the terminal guidance is schematically shown in Fig. 2. DI performed 3 ITMs at Impact $(I)-90$ min, $I-35$ min, and $I-12.5$ min to hit Tempel 1. In DI's case, the comet had been seen for months in the camera, and at the shuttering of the first AutoNav image at $I-2$ h, was about 10 pixels across. In our asteroid scenarios, however, it is quite possible that the asteroid would not have been seen prior to initiation of AutoNav, and is almost always unresolved at this time. We therefore make the assumption that the first ITM is used to correct the bulk of the initial target relative error in the trajectory, which will be a combination of the heliocentric asteroid ephemeris error, and the spacecraft's heliocentric orbit error based on ground radiometric tracking.

Because the size of the maneuver scales as the reciprocal of distance, it is advantageous to perform the first targeting maneuver as early as possible. However, it is also necessary to insure that the target has been positively identified and a good OD solution be available before committing to a maneuver. Thus, the exact timing of the maneuver will be dependent on the particular scenario, but in general, starting AutoNav 2 h before impact is a reasonable starting point. With images taken at a rate of one every 15–30 s to allow for processing time, placing the first ITM at $I-1$ h allows for a good block of data to be accumulated. This has the added benefit of being able to perform outlier checking to make sure bad data, such as a misidentification of spurious signals for the target, that can be rejected. The data is cutoff 2–3 min before the maneuver to allow for the sequence to rotate the spacecraft in the proper orientation (if needed) to perform the burn.

The second burn is needed largely to clean up any errors from the first. Since maneuver execution errors from the first ITM have the longest time to propagate, it is desirable to clean up errors from this maneuver. The timing of ITM 2 is slightly arbitrary and the requirement is to place it somewhere between ITM 1 and ITM 3 such that enough data has accumulated before the OD cutoff to have a good OD fit, and enough time is available after this maneuver to accumulate sufficient data prior to ITM 3. Furthermore, from an operational sense, it is also desirable to perform this maneuver at a time such that: (1) if ITM 1 does not execute successfully, the size of this burn needed to remove ephemeris errors does not use up excessive amounts of fuel, and (2) there is some reasonable chance that if ITM 3 fails to execute, the spacecraft will still impact. For our simulations, we have placed this maneuver at around the $I-30$ min point.

The final maneuver is the most critical as this one will have to guarantee an impact with very high probability. Given the approach velocities and sizes of the asteroids in question, it is likely that the asteroid will have been unresolved in the images through the first two maneuvers. Thus, it is highly desirable that the last maneuver be performed after the asteroid is resolved such that the brightness centroid is unambiguously on the surface for several camera frames. It should be noted, however, that even if the asteroid is resolved, the amount of surface area illuminated will be dependent on the overall shape and aspect ratio of the asteroid, and the approach phase angle. Simulations performed for DI showed that in some cases, the view of the target could be only a thin crescent or even have multiple discrete lobes [12]. The size of the object in the FOV can be calculated by the formula:

$$N_p = d/(V_\infty t IFOV). \tag{20}$$

where N_p is the desired size of the asteroid in pixels, d is the diameter of the target, t is the time to encounter, and V_∞ is the approach velocity. We can also use the formula to solve for the time, picking a value of N_p we want the last image to have; for most of the scenarios we describe below, we chose 5 as a minimum size and do the last burn only after an image with the object at this size is taken. In cases where the V_∞ is very large, however, this is infeasible because the time is too short to process the image, perform the OD, and turn and do the maneuver by the time the target gets to this size. In general, this final maneuver will be performed 2–4 min before the nominal impact.

9. Case study scenarios

As described earlier, AutoNav for DI was optimized carefully for its particular scenario. We would like to broaden the analysis to encompass a range of scenarios that an asteroid impact mission can encounter. In principle, this is a very wide set of conditions, but we can narrow it down by leveraging off other studies on trajectory searches for impact missions. In particular, a recent study

by Hernandez and Barbee performed a search on accessible NEOs for a deflection experiment and narrowed down mission possibilities to them using specific search criteria [3]. Subsequent to that work, they performed an additional search redefining the criteria such that candidate trajectories that minimize the approach phase angle (an important consideration for imaging on approach) were selected (these results have not yet been published). This resulted in a set of 128 trajectories; in Fig. 3 we plot the approach phase angles as a function of the V_∞. The diameters of asteroids represented in this data set range from a minimum of around 100 m and a maximum of near 300 m.

Fig. 3 shows that the V_∞ s can range from about 3 km/s to upwards of 20 km/s, and the phases range from near 0 deg to almost 180 deg. For our simulations, we picked 4 spot locations from the plot that represents a broad range of conditions. These cases are detailed in Table 1 and shown on the plot (for reference, DI's location on this plot is also labeled). Cases 1 and 2 represent relatively benign scenarios at the low end of the approach velocities and low to moderate approach phase angles. Cases 3 and 4 represent more extreme scenarios; Case 3 for a high phase angle and Case 4 for very high approach velocities. We used the size data as a bounds on the target diameters for our scenarios.

For each case listed in the table, two 500 sample Monte Carlo simulations were performed, one with an asteroid with a harmonic mean diameter of 100 m, and the other with 300 m. These represent the lower and higher ends of the asteroid sizes present in the data shown in Fig. 3. For the triaxial ellipsoid model, we assumed an aspect ratio of 1.5:1 for the largest and 2 smaller axes. Thus, the 100 m asteroid has major and minor diameters of 130 and 65 m, while the 300 m asteroids dimensions are 390 and 260 m.

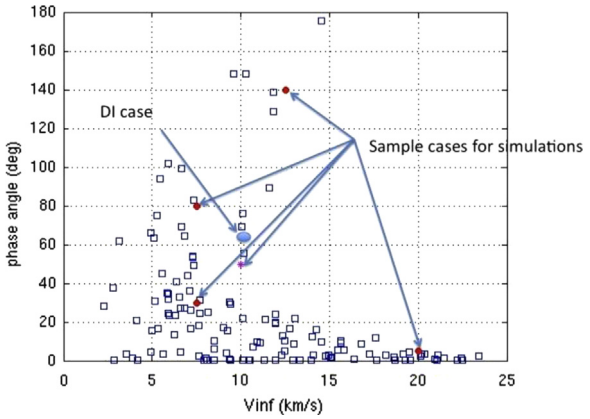

Fig. 3. Phase vs. V_∞ for Candidate Mission Set (data from [3]).

Table 1
Representative scenario cases for simulations.

Case	V_∞ (km/s)	Phase angle (deg)
1	7.5	30
2	7.5	80
3	12.5	140
4	20	5

For the orientation of the asteroid, we randomly sampled a uniform distribution of 0–180 deg in the pole right ascension, and −90 to +90 deg in pole declination (in the Earth Mean Orbital coordinate frame), in each sample of the simulation.

10. Monte Carlo simulations

Linear covariance is generally used to obtain navigation accuracy statistics for deep space missions. This technique, however, has the limitation that all the error sources have Gaussian distributions, which is not always the case. In particular, the error introduced by assuming the center of brightness is the same as the center of mass for an object with an irregular shape, unknown mass distribution, and seen from non-zero phase angles, is not easily modeled in the covariance world. For this reason, we prefer to use Monte Carlo simulations to assess the impactor navigation performance for the terminal guidance phase. The following subsections will detail each of the error sources sampled in the simulations.

10.1. Initial ephemeris error

This is the combination of the spacecraft's and target asteroid's heliocentric ephemeris error, expressed as the error of the spacecraft's cartesian position and velocity relative to the asteroid. Since their respective heliocentric orbits are computed independently (until the asteroid is imaged with the onboard camera which ties them together), the relative error is a root sum square of the two. As mentioned earlier, the asteroid ephemeris error is roughly in the tens of km category; the spacecraft's will be comparable, also in the low tens of km. Thus, in our simulations, we assume an initial target relative orbit error of 30 km, 1σ, in position. The velocity errors are also comparable between the asteroid and the spacecraft, and we sample an uncertainty value of 5 cm/s, 1σ.

10.2. Maneuver execution error

This error accounts for the difference between the desired ΔV for a maneuver and the burn that is actually executed. The causes for the error include sources such as limitations on the accelerometer measurements which cut the burn at the desired ΔV, plume impingement, thruster alignment errors, and attitude errors. Assuming that the spacecraft thrusters are balanced (ideally no net translational ΔV on the spacecraft if thrusters are used to make attitude changes), we also subsume the unbalanced component of turns for the maneuver into the execution error model. The model itself uses the Gates method, which describes the execution error having a fixed and proportional component for the burn value and the pointing [13]. Table 2 lists the values used in the simulations.

10.3. Gyro errors

If attitude knowledge comes from ACS using IMU propagation, then the errors accumulated by the gyros needs to be included. This includes the initial attitude bias

Table 2
Gates maneuver execution error model parameters.

Parameter	Error sampling (1σ)
Fixed magnitude	4.3 mm/s
Proportional magnitude	10%
Fixed direction	4 mm/s
Proportional direction	3.1%

Table 3
Sample gyro errors.

Gyro	Rate bias (1σ) (deg/h)	Angle random walk (deg/$\sqrt{\text{h}}$)
MIMU	0.005	0.005
SSIRU	0.0005	0.0005

caused by the initialization of the absolute attitude from the last star tracker update, plus the gyro bias (which is actually an angle rate) and gyro angle random walk. The attitude bias we assume to be about 150 μrad 1σ as a representative value; for the gyro bias and angle random walk, we take two values for MIMU and SSIRU class gyros. The values for these are given in Table 3.

10.4. Scene generation

In order to get accurate impact statistics, it is important to simulate the image processing as close to reality as possible. Thus, we want to generate photorealistic images for both unresolved and resolved images; the former is important since centroids are obtained in reality to subpixel accuracy. Currently, we have the capability to model the asteroid as a triaxial ellipsoid, and given the dimensions of the ellipsoid and its inertial orientation (randomly selected for each Monte Carlo sample), a realistic scene can be generated. This is done by finding the illumination level of points on the body by tracing the rays from the camera to the object. The equations that determine the illumination are provided in Ref. [14]. Subpixel resolution is obtained by dividing each pixel into smaller sections for tracing the rays, and each of these is convolved with the PSF before summing up the intensities at the pixel level. The amount of subdivision was determined by experimentation; a factor of 2–3 smaller than the angular size of the object was found to be sufficient (Table 1).

10.5. Simulation setup

The simulation setup was as follows. For each Monte Carlo run, the filter is initialized at about $I-2$ h, with a priori values of the state being an ideal trajectory which impacts the center of the comet. The a priori sigmas on the estimated parameters set to be 50 km in position and 10 cm/s in velocity. In the IMU attitude reference mode, the four bias and drift parameters are included in the estimate list with a prior sigmas of 20 pixels and 0.005 pixels/s (MIMU), or 20 pixels and 0.00024 pixels/s (SSIRU). In both cases, the a priori values in the filter for the bias and drift parameters were always set to 0. At 2 min intervals, the scene is generated based on propagation of the "truth" trajectory, which is randomly sampled using the statistics given above for the ephemeris errors. In the IMU mode, attitude knowledge errors accumulated by the IMU are simulated as follows. The image passed to AutoNav is created using attitude information computed by adding the IMU error, randomly sampled from the initial attitude, rate and random walk statistics shown in Table 3, to the nominal attitude assumed by AutoNav. After 30 min, 15 images have been processed by AutoNav and the first OD is performed. Subsequently, an OD solution is done after every image. The OD is cutoff 2 min prior to the maneuver, and the filtered state information is used to predict the miss distance, from which a maneuver is computed. The computed maneuver is used to propagate the spacecraft state in the OD filter; for the "truth" trajectory, maneuver execution errors are added by sampling the statistics in Table 2 and applying the Gates model. After each maneuver, the filter is reinitialized to start at the first image after the maneuver, and the image processing/OD procedure is performed again. We restart the filter with the original a priori uncertainties because we found that carrying over the post-fit covariance overly constrains the filter. In the filter reinitialization, the a priori values for the filter parameters are set to be the last estimate prior to the maneuver, propagated to the current time. Between ITMs 1 and 2, the image frequency is increased to 1 min intervals, and between ITMS 2 and 3, it is increased even further to 30 s intervals. As the truth trajectory either crosses the surface of the triaxial ellipsoid, or it is determined to have passed by the asteroid, the simulation is stopped and the relevant parameters (surface point of impact or closest altitude on flyby, magnitudes of maneuvers, OD errors) stored. The process is then repeated for as many Monte Carlo samples as needed. The camera parameters used for all the simulations were the same as that on the DI impactor, with an IFOV of 10 μrad.

11. Monte Carlo results

11.1. Preliminary results using subset of full simulation

Before describing the results of the full set of simulations, we did a preliminary Monte Carlo analysis to determine, in a general sense, what the effect of the most critical parameter, the attitude knowledge error, is to the OD solution and hence, targeting accuracy. For this set of simulations, instead of creating simulated images, we simply used the ideal center of the target in the image as the observation to eliminate asteroid size and phase effects, and set the maneuver execution error to 0. We used the following scenario as a reasonable representative:

- Asteroid approach velocity of 10 km/s
- Camera with 10 μrad IFOV
- Maneuvers at $I-1$ h, $I-30$ min, and $I-5$ min

Sampling the initial spacecraft target relative state and attitude errors, we performed a Monte Carlo simulation with attitude knowledge errors of 0 (stellar mode), MIMU

level, and SSIRU level. We then collected the information about the miss distance in the asteroid-centered B-plane, plotting the cumulative probability of miss distances at increments starting from 10 m on up. The result is shown in Fig. 4 for the three attitude knowledge cases. From this, it can be seen that in the stellar reference case, 100% of the cases are reached at a miss distance of 50 m; for the SSIRU, this is reached by 100 m, and for the MIMU, it occurs only at about 800 m. Thus, we can infer that if the target asteroid is less than about 800 m, the project will have to commit to either spending money on the higher quality gyro, or find a combination of camera and approach asymptote that allows for visible stars up to the final image. Conversely, if stellar reference is available during approach, the size of the asteroid can be fairly small and still have a reasonable chance of impacting. Based on this result, for the subsequent analysis we did not include results for the MIMU gyros since all our sample asteroids were less than 300 m in diameter.

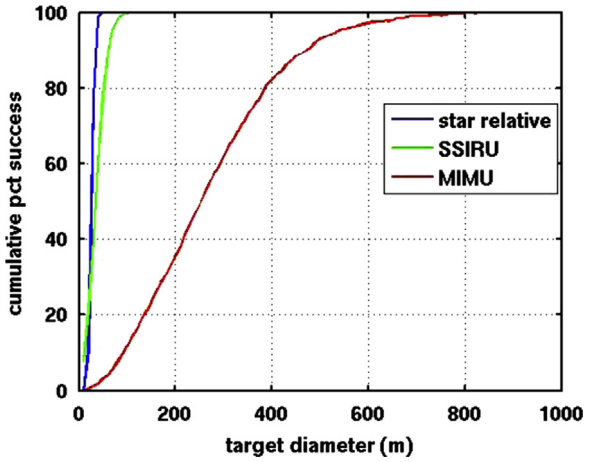

Fig. 4. Cumulative probability as a function of miss distance for three attitude modes for the preliminary Monte Carlo analysis.

Table 4
Probability of asteroid impact.

Case	Stellar reference		SSIRU	
	100 m (%)	300 m (%)	100 m (%)	300 m (%)
1	98.8	100.0	85.5	100.0
2	96.5	100.0	73.8	99.2
3	56.6	99.4	53.8	90.6
4	100.0	100.0	75.4	99.6

11.2. Results using full simulation

Table 4 shows the results of the full simulation (including simulated images and maneuver execution errors) for each of the cases for the stellar reference and SSIRU attitude knowledge cases, for the 100 m asteroid and the 300 m asteroid. The results are given in terms of the percentage of time the spacecraft impacted the asteroid for each scenario. Immediately obvious from this table is the sensitivity to both attitude reference and the phase angle. Cases 2 and 3, which have approach phases of 80 and 140 deg, respectively, have noticeably lower chances of success than the others. For Case 3 in particular, the odds of hitting a 100 m asteroid are barely above 50%, thus pointing to a scenario which should be avoided if possible. The reason the high phase is a problem can be seen from Fig. 5, which shows a zoomed-in view of what the asteroid looks like at the final image taken before ITM3 at the phase angles represented by Cases 2, 3, and 4. Also shown on the plots is the computed COB (o), and the true center of the object (+). At 5 deg, the two are almost on top of each other, but for 140 deg, the COB, due to the PSF, is very near the edge of the physical object.

One somewhat surprising result of this analysis was that the very high velocity case had relatively few misses, except for the 100 m asteroid with SSIRU attitude reference. In this scenario, the last image was taken at $I-2$ min, where the asteroid was roughly 4 pixels across – not enough to clearly see the shape. Nevertheless, due to the fact that the centroid is so close to the true object center as seen in Fig. 5(a), the OD was accurate enough to do precise targeting. This result is promising, as it means that deflections can be increased with higher V_∞ without sacrificing targeting accuracy.

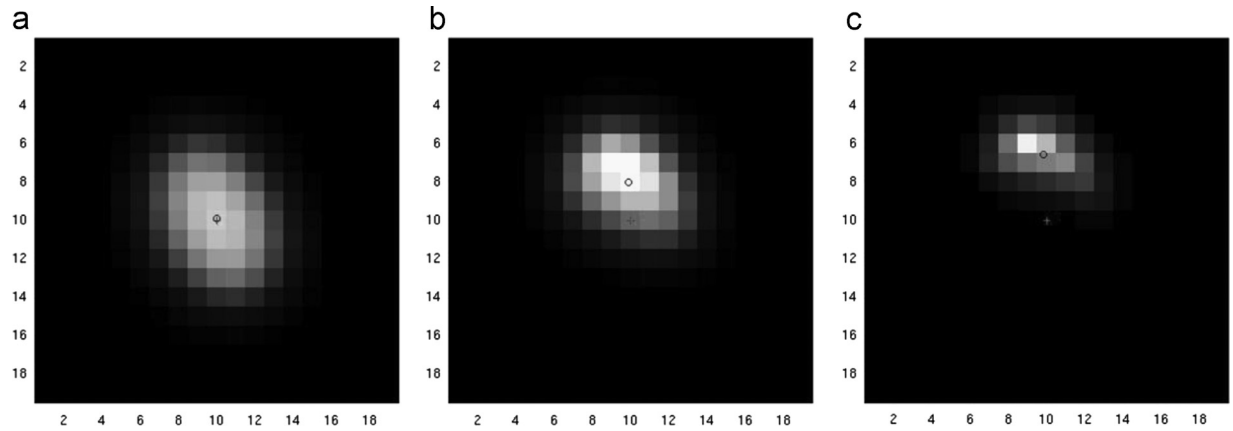

Fig. 5. Final image before ITM3 for 3 phase angles. (a) Phase=5 deg. (b) Phase=80 deg. (c) Phase=140 deg.

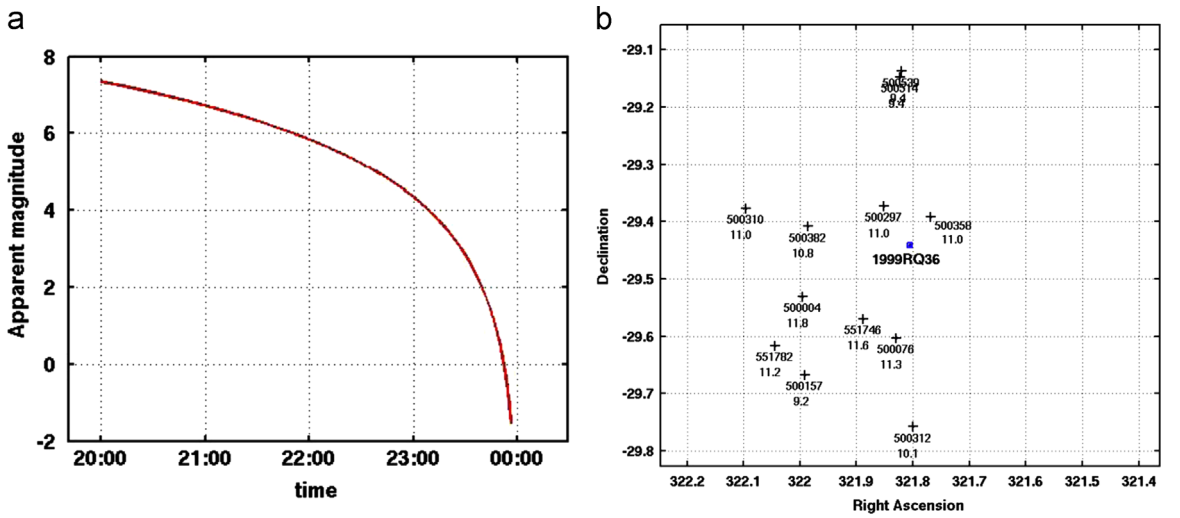

Fig. 6. Apparent magnitudes of 1996RQ36 and background stars during terminal guidance phase. (a) Apparent magnitude of 1996RQ36 on approach. (b) Background stars on approach.

11.3. ISIS example

In addition to the generic cases described above, we also analyzed a specific example case – the ISIS mission. ISIS (Impactor Surface and Interior Science) is a proposed mission to send an impactor spacecraft to hit the asteroid 1999RQ36 in early February 2021 [15]. This particular asteroid was chosen because NASA's OSIRIS-REx spacecraft will be resident there during this time, having obtained its sample and waiting for the correct planetary alignment to begin the return journey to Earth. Thus, OSIRIS-REx will be in a unique position to observe the ISIS impact with the goal of obtaining scientific data on the resultant crater and measure the deflection caused by the impact.

Although several trajectory options for ISIS are available, the preferred one has an approach V_∞ of 13.4 km/s and an approach phase of 9 deg. Due to recent radar measurements, the physical characteristics of 1999RQ36 are fairly well known. We used triaxial ellipsoid dimensions of $517 \times 500 \times 460$ m and a rotation rate of 4.2 h. The pole is known to be almost perpendicular to the ecliptic, so we defined the pole declination to be -90 deg. For the Monte Carlo simulations, we used the same scenario as described above, with the first two ITMs at $I-1$ h and $I-30$ min as before, but ITM 3 was placed at $I-3$ min. The DI MRI camera was used for the approach imaging.

For determining the attitude mode, since a reference trajectory is available, we can compare the brightness of the asteroid during the terminal guidance phase against the star background. In Fig. 6, the left panel shows the apparent magnitude of 1999RQ36; in the right panel, the location of the asteroid in the camera FOV is shown against the star background. It can be seen in these figures that the asteroid brightness varies from a minimum of about magnitude 7.5 to a maximum of almost -2. In contrast, the stars in the FOV are all dimmer than 9. Thus, if stellar reference mode is to be used, alternate frames must be taken. We took the more conservative approach and

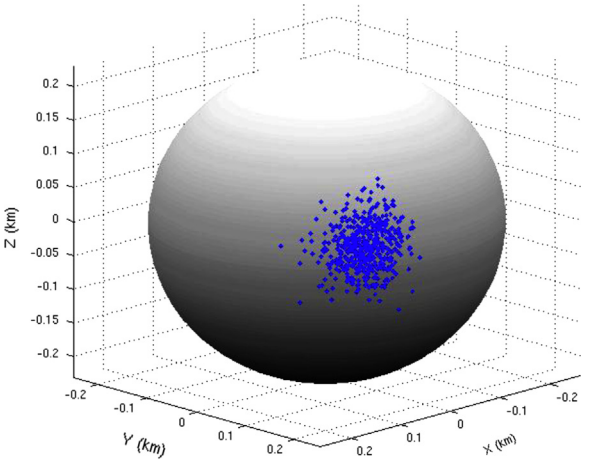

Fig. 7. Impact locations on surface of 1999RQ36.

assumed that the terminal guidance is done using the SSIRU, and thus used the values in Table 3.

The simulation result from a 500 sample Monte Carlo run showed a 100% success rate of impact. The scatter of impact points on the surface is shown in Fig. 7, where it can be seen that all of the impact points are within roughly 100 m radius surrounding the nominal impact point through the exact center of figure of the asteroid.

12. Conclusions

This paper expanded on the experience with high velocity impactor missions obtained from DI and quantified targeting accuracies for a variety of scenarios. It was shown that the attitude reference used and approach phase angles were the primary drivers of the accuracies. The results were shown for asteroid sizes almost two orders of magnitude smaller than the DI target. With the accuracies obtained,

reasonably high confidence of achieving mission success for most of the scenarios that will be faced.

The results in this paper assume spacecraft hardware with similar characteristics to that of DI. This is based on the assumption that, in the current budget environment, any deflection demonstration mission will have to be low-cost. This in turn implies that the hardware must have high heritage in deep space applications, for which DI provides the perfect example. In the case of an actual deflection where success is an absolute necessity, using a previously successful approach has high value to minimize risk. Our results show where this is applicable, and in the cases where the odds are low of impacting, point to where improvements in hardware or techniques should be focused. For example, the need for stable attitude reference indicates hardware configurations using combinations of multiple redundant, high resolution cameras could be used for scenarios with small asteroids. For very high phase angles, the use of visible imagers is necessarily limited and other types, such as infrared detectors, may be required. In all cases, better OD/filtering strategies would be helpful, and future work will focus on this aspect, as well as increasing the fidelity of various aspects of the simulation, including using more unusual shape models, and incorporating finite duration burns.

Acknowledgments

The authors would like to thank Ed Riedel at JPL for creating the code used to create photorealistic images, and Brent Barbee at Goddard Space Flight Center and Sonia Hernandez at the University of Texas for providing the data on mission scenarios.

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

[1] D. Yeomans, Near-Earth Objects: Finding Them Before They Find Us, 1st ed. Princeton University Press, Princeton, NJ, 2012.
[2] D. Yeomans, P. Chodas, M. Keesey, S. Ostro, J. Chandler, I. Shapiro, Asteroids and comets using radar data, Astron. J. 103 (1992) 303–317.
[3] S. Hernandez, B.W. Barbee, Design of spacecraft missions to test kinetic impact for asteroid deflection, in: AAS/AIAA Space Flight Mechanics Conference, Paper AAS 12-129, 2012.
[4] L. Wood, The evolution of deep space navigation: 1962–1989, in: AAS Guidance and Control Conference, Paper AAS 08-051, 2008.
[5] S. Bhaskaran, M. Abrahamson, S. Chesley, M. Chung, A. Halsell, R. Haw, C. Helfrich, D. Jefferson, B. Kennedy, T. McElrath, W. Owen, B. Rush, J. Smith, T. Wang, C. Yen, Navigation of the epoxi spacecraft to comet Hartley 2, in: AAS/AIAA Astrodynamics Specialist Conference, Paper AAS 11-486, 2011.
[6] S. Bhaskaran, Autonomous navigation for deep space missions, in: SpaceOps 2012 Conference, Stockholm, Sweden, 2012.
[7] J. Riedel, T. Wang, R. Werner, A. Vaughan, D. Myers, N. Mastrodemos, G. Huntington, C. Grasso, R. Gaskell, D. Bayard, Configuring the deep impact autonav system for lunar, comet and mars landing, in: AIAA/AAS Astrodynamics Specialist Conference, Paper AIAA 2008-6940, 2008.
[8] D. Kubitschek, N. Mastrodemos, R. Werner, B. Kennedy, S. Synnott, G. Null, S. Bhaskaran, J. Riedel, A. Vaughan, Deep impact autonomous navigation: the trials of targeting the unknown, in: AAS Guidance and Control Conference, Paper AAS 06-081, 2006.
[9] W. Owen, Methods of optical navigation, in: AAS/AIAA Space Flight Mechanics Conference, Paper AAS 11-215, 2011.
[10] M. Abrahamson, B. Kennedy, S. Bhaskaran, Autonav design and performance for the Epoxi Hartley 2 flyby, in: SpaceOps 2012 Conference, Stockholm, Sweden, 2012.
[11] B.D. Tapley, B.E. Schutz, G.H. Born, Statistical Orbit Determination, 1st ed. Elsevier Academic Press, San Diego, CA, 2004.
[12] N. Mastrodemos, D. Kubitschek, R. Werner, B. Kennedy, S. Synnott, G. Null, J.E. Riedel, S. Bhaskaran, A. Vaughan, Autonomous navigation for Deep Impact, in: AAS Guidance and Control Conference, Paper AAS 06-177, 2006.
[13] C.R. Gates, A simplified model of midcourse maneuver execution errors, JPL-Technical Report 32-504, October 15, 1963.
[14] J. Riedel, W. Owen, J. Stuve, S. Synnott, R. Vaughan, Optical navigation during the Voyager Neptune encounter, in: AIAA/AAS Astrodynamics Conference, Paper AIAA 90-2877, 1990.
[15] S. Chesley, J. Elliot, S. Bhaskaran, T. Lam, P. Abell, E. Asphaug, D. Lauretta, The ISIS mission concept: an impactor for surface and interior science, in: Planetary Defense Conference 2013, IAA-PDC2013-04-01, 2013.